

EV369763782

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Methods and Systems for Hinting Fonts

Inventor(s):

David H. Salesin
Geraldine Wade
Douglas E. Zongker

ATTORNEY'S DOCKET NO. ms1-539usc1

1 **RELATED APPLICATION**

2 This application is a continuation of and claims priority to U.S. Patent
3 Application No. 09/620,618, filed on 07/21/2000, the disclosure of which is
4 incorporated by reference.

5

6 **TECHNICAL FIELD**

7 This invention relates to methods and systems for hinting fonts, particularly
8 TrueType fonts.

9

10 **BACKGROUND**

11 The demand for high-quality hinted fonts is outstripping the ability of
12 digital typography houses to produce them. Hinting is a painstaking manual
13 process that can only be done well by a handful of highly skilled professionals. It
14 requires a blend of typographical artistry with technological ability. In order to
15 provide a full appreciation of the hinting problem, a review of how digital fonts
16 are scan-converted onto a raster display is given.

17 In digital typography, each character in a font is described by a set of
18 outlines, usually represented by splines. When the character is rendered onto a
19 grid of pixels, the outlines are scaled to the desired size, and then each pixel whose
20 center lies inside of an outline is set to black. When fonts are displayed at
21 sufficiently high resolutions, this approach works beautifully. But for sizes below
22 about 150 *ppem*, severe aliasing problems can result when this naive outline filling
23 process is applied, especially for delicate features such as serifs. As an aside,
24 hinters express font sizes in *pixels per em*, or *ppem*. This measure counts the
25 number of device pixels in the *em* of the font. In traditional typography, the *em* of

1 a font was the height of the metal block of type. With digital typography, there is
2 no actual metal block to measure, but the digital outlines are still expressed in
3 coordinates relative to this hypothetical size. The point size of text refers to the
4 size of its *em* expressed in points (a point is 1/72 of an inch). Thus, “12-point
5 text” would correspond to 12 ppm on a 72 dpi screen, or 100 ppm on a 600 dpi
6 printer.

7 Fig. 1 shows an example of the aliasing problems that can result when a
8 naïve outline filling process is applied. Consider the leftmost image 100 in the
9 form of a lower case letter “a” which represents an outline of the character that is
10 to be filled. The centermost image 102 (absent the illustrated black pixels 102a) is
11 generated by a naïve algorithm. This pixel pattern does not look much like a
12 lowercase “a”. A simple dropout control mechanism can be added to the fill
13 algorithm to turn on additional pixels to preserve the character’s topology. The
14 resultant pixels that are turned on by the dropout control mechanism are shown in
15 black at 102a. The rightmost image 104 illustrates the work of an experienced
16 hinter. It will be observed that the pixel pattern has been subtly altered to both
17 improve readability and better preserve the character of the original outline.

18 The hinting process is not just about optimizing individual characters. The
19 hinter must balance the needs of a single glyph with the desire for consistency
20 across all the characters of a font. It is important, for example, to ensure that all
21 the vertical stems of a font are the same number of pixels wide at a given size. If
22 the scaling and rounding process produces one-pixel-wide stems on some
23 characters and two-pixel-wide stems on others, then a passage of text will look
24 blotchy and be difficult to read. The goal of the hinter is to produce a smooth
25 transition from very high sizes, where merely filling the outlines suffices and

1 hinting is unnecessary, down to lower sizes, where legibility must be preserved
2 even when that means a departure from the outlines drawn by the original font
3 designer.

4 Although the ever-improving resolution of hardcopy devices is beginning to
5 approach the point at which hinting is not necessary, the technology is not there
6 yet: 10- or 12-point text on a 300 or even 600 dpi printer still needs hinting for
7 best results. More importantly, the increasing emphasis on reading text on-
8 screen—from visions of the “paperless office” to the emergence and proliferation
9 of hand-held computers and eBooks—means that more and more text is being
10 viewed on devices in the 72–100 dpi range. Though resolutions of these displays
11 are improving as well, for the foreseeable future hinting will be an absolute
12 necessity in order to provide clear, legible text.

13 Although attempts have been made to design automated hinting systems in
14 the past, even the best of these produce hints that are good, but still not up to the
15 standards of professional typographers. Exemplary systems are described in the
16 following references: Andler, *Automatic generation of gridfitting hints for*
17 *rasterization of outline fonts or graphics*, Proceedings of the International
18 Conference on Electronic Publishing, Document Manipulation, and Typography,
19 pps. 221-234, Sept. 1990, and Hersch, *Character generation under grid*
20 *constraints*, Proceedings of SIGGRAPH 87, pps. 243-252, July 1987.

21 This previous work assumed that in order to be useful, an autohinter had to
22 be a monolithic, self-contained package: outlines in, quality hints out. That is an
23 admirable goal, and it may be achieved someday. However, given the detailed,
24 aesthetically-based nature of the work, a better, more useful approach is to view
25 the autohinter as one piece of a system that includes a human hinter.

1 There are two major font standards in widespread use today: Type 1 and
2 TrueType. Type 1 fonts (Adobe Systems Inc., *Adobe Type 1 Font Format*, March
3 1990), often called “PostScript fonts,” were developed by Adobe and are popular
4 in the world of publishing. Printing applications were the target when this system
5 was developed, though utilities are now available to enable on-screen display of
6 Type 1 fonts. The TrueType format (Apple Computer, Inc., *The TrueType Font*
7 *Format Specification*, 1990, Version 1.0), originally developed by Apple, was
8 intended to unify type on the screen and on paper, and is used in both the
9 Macintosh and Windows operating systems. TrueType has something of a
10 reputation for being of low quality, but this is mostly due to the fact that TrueType
11 was always an open standard while Type 1 was not, and so the public domain is
12 flooded with a large number of poorly designed, unhinted TrueType fonts. The
13 TrueType standard does contain extensive facilities for high-quality hinting,
14 though, and through the efforts of Monotype and other typography houses, more
15 and more quality fonts are now available in the TrueType format.

16 Though both formats represent characters as spline-based outlines, the
17 hinting styles are radically different. Hinting for Type 1 fonts works by marking
18 sections of the outline as corresponding to particular typographic features of the
19 character—stems, bowls, counters, and so on. It is the job of the rasterizer to take
20 advantage of these hints about the character shape to produce the best possible
21 pattern of pixels. This scheme has the advantage that enhancements to the
22 rasterizer can produce improvements to all fonts on the system, but means that a
23 designer of digital type cannot specify exactly what an outline will look like when
24 rendered at a given size.

1 The TrueType font technology takes a different approach. Instead of
2 leaving control over the glyph's final appearance to the rasterizer, a TrueType font
3 contains explicit instructions (hints) about how particular control points should be
4 shifted to fit the pixel grid. These instructions take the form of a program in a
5 special, TrueType-specific bytecode language. Since both the behavior of each
6 instruction and the rasterizing algorithm are defined in the TrueType standard, the
7 designer of a TrueType font can predict exactly which pixels will be turned on for
8 a character at a given size, no matter what type of output device is being used.

9 In TrueType, each contour of an outline is specified with a sequence of
10 point positions (also referred to herein as "control points" or "knots"). Fig. 2
11 shows exemplary outline curves of the lowercase letter "a" for two fonts:
12 Trebuchet and Frutiger. Each point is flagged as either "on-curve" or "off-curve".
13 TrueType defines the outline as follows:

- 14 • Two successive on-curve points are connected with a straight line
15 segment.
- 16 • When an off-curve point falls between two on-curve points, the three
17 are treated as the control points for a quadratic Bézier segment.
- 18 • When two adjacent off-curve points appear, the midpoint of the
19 segment connecting them is treated as an implicit on-curve point
20 between them, allowing reduction to the case above.

21 The glyph renderer starts by scaling the outlines to a particular size, then
22 executes the attached program to shift control points around in a size-specific way
23 before filling the altered outline. By itself, this approach cannot produce the
24 necessary consistency among different characters of a font, or even between
25 different parts of the same character, since each action is necessarily local. Global
synchronization of outline alterations is achieved through use of a control value

1 table, or CVT. The CVT is a shared table of distances, that can be referenced by
2 instructions in each glyph's program. When the rendering is initialized for a given
3 size, the values in the CVT are scaled and rounded to the current grid size. Point
4 movements can then be constrained by CVT entries. For instance, a person
5 writing hints for TrueType may decide to use CVT 81 to represent, say, the width
6 of vertical black stems in lowercase letters. He or she will then write instruction
7 sequences for all appropriate lowercase letters, all referring to CVT entry 81, so
8 that all the vertical black stems at a given size will have the same width.

9 The TrueType language is an assembly-style stack-based language. The
10 intent of the designers of TrueType was not to make typographers learn and write
11 in the TrueType language itself, but rather to facilitate the development of high-
12 level languages and tools that generate TrueType code. The Visual TrueType
13 (VTT) package from Microsoft is such a tool. VTT is described in detail in
14 Stamm, *A Graphical Method for Authoring Font Intelligence*, Electronic
15 Publishing, Raster Imaging, and Digital Typography, pps. 77-92, March/April
16 1998.

17 VTT provides a high-level language, called "VTT Talk", for expressing
18 relationships between points. VTT Talk provides statements for expressing the
19 following classes of hints:

- 20 • *Link* constraints: the vertical or horizontal distance between a pair of
21 knots is constrained by an entry in the CVT.
- 22 • *Dist* constraints: the "natural" vertical or horizontal distance between
23 a pair of knots is maintained, so that if one point is moved the other
moves in parallel.
- 24 • *Interpolate* constraints: a knot's fractional distance between two
parent knots is maintained.

1

- *Anchors*: specific knots can be rounded to the nearest gridline, or to
2 a gridline specified by a CVT entry.

3

4 These types of hints are demonstrated visually for two characters from the
5 Georgia Roman font in Fig. 3. There, the constraints are labeled for identification.

6 The VTT Talk hints are compiled into a TrueType program stored in the
7 font file. One advantage of working with VTT Talk is that each statement simply
8 asserts a relationship between two points, and there is little dependence on the
9 order of the statements. If one statement is omitted, the meaning of the others is
10 unchanged. In contrast, TrueType assembler is a sequential language that
11 maintains a fairly complex state. Most instructions in TrueType have side effects
12 that modify this state. If one tried to translate the assembler code directly, and
13 were for some reason unable to translate a particular instruction—for instance, due
14 to a sufficiently large difference in the matched glyphs' outlines—the effects of
subsequent instructions could change entirely.

15 This invention arose out of concerns associated with improving the systems
16 and methods through which hinting takes place. In particular, the invention arose
17 out of concerns associated with improving TrueType hinting systems and tools.

18

19 **SUMMARY**

20 Methods and systems for hinting fonts, particularly TrueType fonts, are
21 described. The inventive methods and systems provide a mechanism by which
22 hints are automatically translated from one font to another. An advantageous
23 feature of the described approach is that it preserves the basic strategy and
24 structure of the original hints, which were handcrafted by a professional
25 typographer for each individual glyph of the font. Generally, these translated hints

1 provide an excellent starting point for a human typographer to fine-tune and ad-
2 just.

3 In the illustrated and described embodiment, a character or glyph (i.e. a
4 source character) from a first font is selected and provides hints that are to be
5 transferred to a character or glyph of a second font (i.e. a target character). The
6 hints, in the described embodiment, are TrueType hints that are statements defined
7 in terms of control points or knots that define the shape or appearance of a
8 character. The hints constrain the character's control points, e.g. the distance
9 between a pair of control points or how the points relate to one another. Outlines
10 or contours for the two characters are first matched to establish a correspondence
11 between the characters. Next, various attempts are made to find a "best match"
12 between individual control points on the source character and target character.

13 In the illustrated and described embodiment, finding a best match between
14 individual control points on the different characters involves a process by which
15 control points on the source character are paired with different control points on
16 the target character. Each pairing of control points is then scored in accordance
17 with predefined criteria to provide a local feature score for each control point pair.
18 The local feature score is a measure of how well the individual control points
19 match. The local feature scores are then summed to provide an overall score of
20 the set of paired control points. Different control point pairings are tried and result
21 in multiple sets of paired control points. Each of the sets of paired control points
22 has an overall score.

23 A predetermined number of the best overall scores are selected and then
24 each overall score is attempted to be improved upon. This takes place by
25 manipulating the individual control point pairings within each set that corresponds

1 to an overall score, subject to a condition designed to preserve the order of the
2 control points. The result of this process is a single score that reflects the best or
3 most optimal pairing of control points.

4 Once a desirable pairing of control points is ascertained, hints from the
5 source character can be translated to the target character. Advantageously, in the
6 described embodiment, translation of the hints involves manipulating control point
7 references in the defined hints so that the hints now refer to control points in the
8 target character rather than control points in the source character. Because the
9 most optimal pairing of control points was ascertained, this process is greatly
10 facilitated.

11 In addition, the hints can contain references to a control value table (CVT).
12 The CVT contains table entries that are associated with values that are used to
13 constrain the control points of the source character. In accordance with one
14 embodiment, CVT values are manipulated so that the entries now correspond to
15 values that are used to constrain the character in the target font. Manipulating the
16 CVT values recognizes that there are differences between the source and target
17 characters and that using the values associated with the source character in
18 connection with the target character would be inappropriate.

19

20 **BRIEF DESCRIPTION OF THE DRAWINGS**

21 Fig. 1 illustrates outline for the Palatino Italic “a”, along with a pixel
22 pattern generated by rasterizing the outlines for display of 18-point text on a 72
23 dpi device.

24 Fig. 2 illustrates, for two different TrueType fonts, various “on-curve” and
25 “off-curve” control points that are used to define hints for the fonts.

1 Fig. 3 illustrates a visualization of VTT Talk hints created by a professional
2 hinter for two characters of Georgia Roman.

3 Fig. 4 is a high level diagram of a computer system that can be utilized to
4 implement an autohinter in accordance with the described embodiment.

5 Fig. 5 is a flow diagram that describes steps in a method in accordance with
6 the described embodiment.

7 Figs. 6-8 illustrate different features that are utilized to ascertain a score for
8 a pair of control points in accordance with the described embodiment.

9 Fig. 9 illustrates one way in which control points can be paired for purposes
10 of generating a score.

11 Fig. 10 illustrates a final match between control points for the two
12 characters of Fig. 2 in accordance with the described embodiment.

13 Fig. 11 illustrates a final match for two different characters in accordance
14 with the described embodiment.

15 Fig. 12 is a flow diagram that describes steps in a method in accordance
16 with the described embodiment.

17 Fig. 13 is a flow diagram that describes steps in a method in accordance
18 with the described embodiment.

19 Fig. 14 visually illustrates the outcome of a process in which control value
20 table (CVT) values for one font were transferred to another font.

21 Figs. 15-19 visually illustrate hints transferred from one font to another.

22 Fig. 20 illustrates the Sylfaen Sans Bold font, both unhinted and with hints
23 transferred from Sylfaen Sans Roman.

24 Fig. 21 illustrates the Georgia Bold font, both unhinted and with hints
25 transferred from Georgia Roman.

1

2 **DETAILED DESCRIPTION**

3 **Overview**

4 The inventive approach described below was primarily motivated by the
5 work of Hersch and Betrisey, see, *e.g.* Hersch and Betrisey, *Model-based*
6 *Matching and Hinting of Fonts*, Proceedings of SIGGRAPH 91, pps. 71-80, July
7 1991. In that method, hints are generated for each glyph by matching its outline to
8 a human-constructed generic model of that character's shape (for example, a
9 generic uppercase roman 'B'). The model consists of two representations of the
10 generic character shape. A *skeleton model* builds the character out of solid parts,
11 labeled as stems, bowls, serifs, and so on. A *contour model* is an outline
12 representation of the character, constructed to have as few control points as
13 possible while still spanning the space of possible character shapes. The
14 correspondences between the two models are known because they are specified by
15 hand when the model is built. In their method, the outlines of the glyph to be
16 hinted are matched to the corresponding contour model by a fairly complex
17 process that takes into account both global and local features. Points are classified
18 by their position relative to the baseline, cap-height and x-height lines, and left and
19 right sidebearings. Local features distinguishing points are based on the curvature,
20 direction, and orientation of the adjacent curve segments. Once the
21 correspondence between the unknown outline and the model outline is established,
22 the known correspondence between the model outline and the model skeleton can
23 be used to label parts of the unknown outline as belonging to significant features
24 such as stems and serifs. From this labeling a set of Type 1-style hints for the new
25 outline can be derived.

1 Hersch and Betrisey's work requires a manually constructed model in order
2 to link points on the outline with the "semantic" features needed for hinting.
3 Hinting in TrueType does not require an explicit labeling of these features. Rather,
4 this information is implicitly used by the human typographer when deciding on a
5 hinting strategy for the character. The end result expressed in the font is simply a
6 set of relationships, or constraints, between control points. These constraints
7 obviate the need for the skeleton model—once we find the correspondence
8 between a contour model and the outlines of the target glyph, hints can be
9 immediately produced for the target outline without transitively applying a second
10 correspondence.

11 In the inventive approach, we have reduced our needs to having a contour
12 model with control-point-level hints attached to it. Advantageously, a TrueType
13 font that has already been hinted is used as the contour model. This has a number
14 of advantages over using a specialized model built expressly for the auto-hinter.

15 First, there are immediately a wide variety of fonts from which to choose as
16 templates. Moreover, choosing a template close to the target font will increase the
17 likelihood of a good match, and consequently the quality of the resulting hints.
18 Advantageously, a template font can be selected from a library automatically, or
19 within a typeface classification system, or different template fonts can be chosen
20 for different characters of the target. Another advantage of using real hinted fonts
21 as templates is that typographers, rather than computer scientists, can build
22 templates using tools with which they familiar. Furthermore, each typographer
23 can build templates to suit his or her hinting style.

24
25

Exemplary Computer System

Fig. 4 is a high level block diagram of an exemplary computer system 130 that can be programmed to function as an automated hinting system referred to as an “autohinter”.

Computer 130 includes one or more processors or processing units 132, a system memory 134, and a bus 136 that couples various system components including the system memory 134 to processors 132. The bus 136 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. The system memory 134 includes read only memory (ROM) 138 and random access memory (RAM) 140. A basic input/output system (BIOS) 142, containing the basic routines that help to transfer information between elements within computer 130, such as during start-up, is stored in ROM 138.

Computer 130 further includes a hard disk drive 144 for reading from and writing to a hard disk (not shown), a magnetic disk drive 146 for reading from and writing to a removable magnetic disk 148, and an optical disk drive 150 for reading from or writing to a removable optical disk 152 such as a CD ROM or other optical media. The hard disk drive 144, magnetic disk drive 146, and optical disk drive 150 are connected to the bus 136 by an SCSI interface 154 or some other appropriate interface. The drives and their associated computer-readable media provide nonvolatile storage of computer-readable instructions, data structures, program modules and other data for computer 130. Although the exemplary environment described herein employs a hard disk, a removable magnetic disk 148 and a removable optical disk 152, it should be appreciated by

1 those skilled in the art that other types of computer-readable media which can
2 store data that is accessible by a computer, such as magnetic cassettes, flash
3 memory cards, digital video disks, random access memories (RAMs), read only
4 memories (ROMs), and the like, may also be used in the exemplary operating
5 environment.

6 A number of program modules may be stored on the hard disk 144,
7 magnetic disk 148, optical disk 152, ROM 138, or RAM 140, including an
8 operating system 158, one or more application programs 160, other program
9 modules 162, and program data 164. A user may enter commands and information
10 into computer 130 through input devices such as a keyboard 166 and a pointing
11 device 168. Other input devices (not shown) may include a microphone, joystick,
12 game pad, satellite dish, scanner, or the like. These and other input devices are
13 connected to the processing unit 132 through an interface 170 that is coupled to
14 the bus 136. A monitor 172 or other type of display device is also connected to the
15 bus 136 via an interface, such as a video adapter 174. In addition to the monitor,
16 personal computers typically include other peripheral output devices (not shown)
17 such as speakers and printers.

18 Computer 130 commonly operates in a networked environment using
19 logical connections to one or more remote computers, such as a remote computer
20 176. The remote computer 176 may be another personal computer, a server, a
21 router, a network PC, a peer device or other common network node, and typically
22 includes many or all of the elements described above relative to computer 130,
23 although only a memory storage device 178 has been illustrated in Fig. 4. The
24 logical connections depicted in Fig. 4 include a local area network (LAN) 180 and
25 a wide area network (WAN) 182. Such networking environments are

1 commonplace in offices, enterprise-wide computer networks, intranets, and the
2 Internet.

3 When used in a LAN networking environment, computer 130 is connected
4 to the local network 180 through a network interface or adapter 184. When used
5 in a WAN networking environment, computer 130 typically includes a modem 186
6 or other means for establishing communications over the wide area network 182,
7 such as the Internet. The modem 186, which may be internal or external, is
8 connected to the bus 136 via a serial port interface 156. In a networked
9 environment, program modules depicted relative to the personal computer 130, or
10 portions thereof, may be stored in the remote memory storage device. It will be
11 appreciated that the network connections shown are exemplary and other means of
12 establishing a communications link between the computers may be used.

13 Generally, the data processors of computer 130 are programmed by means
14 of instructions stored at different times in the various computer-readable storage
15 media of the computer. Programs and operating systems are typically distributed,
16 for example, on floppy disks or CD-ROMs. From there, they are installed or
17 loaded into the secondary memory of a computer. At execution, they are loaded at
18 least partially into the computer's primary electronic memory. The invention
19 described herein includes these and other various types of computer-readable
20 storage media when such media contain instructions or programs for implementing
21 the steps described below in conjunction with a microprocessor or other data
22 processor. The invention also includes the computer itself when programmed
23 according to the methods and techniques described below.

24 For purposes of illustration, programs and other executable program
25 components such as the operating system are illustrated herein as discrete blocks,

1 although it is recognized that such programs and components reside at various
2 times in different storage components of the computer, and are executed by the
3 data processor(s) of the computer.

4

5 **Exemplary Processing Technique Overview**

6 Fig. 5 shows a flow diagram that describes steps in a processing method in
7 accordance with the described embodiment. The method is advantageously
8 implemented in software in connection with an automated programmed autohinter,
9 although it could be implemented in any suitable hardware, software, firmware, or
10 combination thereof.

11 Step 500 provides a source character from which hints are going to be
12 translated or transferred. As was pointed out above, the source character is
13 advantageously a character that has already been hinted. This constitutes a
14 dramatic improvement in flexibility and processing time over other methods in
15 which hints are not translated from previously-hinted fonts. Step 502 provides a
16 target character to which the hints from the source character are to be translated or
17 transferred. Step 504 matches contours or outlines of the source and target
18 characters. The goal of this step is to identify and pair contours on the source and
19 target characters that are like contours. For example, the lower case letter "a" has
20 two contours—an inner contour (defining the enclosed inner area in the body of
21 the letter) and an outer contour (defining the outer perimeter and hence shape of
22 the letter). This step ensures that like contours from the source character and the
23 target character are paired together. A specific example of one way of
24 accomplishing this task is described in the section entitled "Matching the
25 Outlines" below.

1 Once the contours are matched, step 506 attempts to match knots (i.e.
2 control points) on one of the contours with knots on the other of the contours.
3 Throughout this document, the terms “knots” and “control points” will be used
4 interchangably. Recall that each contour is defined by or associated with a number
5 of control points that are either on-curve or off curve. What this step attempts to
6 do is to find a “best match” between the control points on the matched contour
7 pair. In the illustrated and described embodiment, this step is accomplished by
8 establishing quantifiable criteria and then examining and scoring knot pairs on the
9 two contours. It will be appreciated that step 506 is performed a number of
10 different times for each pair of matched contours. This produces, in turn, a
11 quantified measurement that gives an indication of the desirability of a particular
12 association of knots between the outlines. From these quantified measurements, a
13 best measurement is selected as constituting the “best match” between the control
14 points on the different outlines or contours. An example of one way of
15 accomplishing this task is given below in the section entitled “Calculating a
16 Score”.

17 Step 508 then checks to see whether there are any additional contour pairs
18 for evaluation. If there are, step 510 gets the next pair of contours and returns to
19 step 506. If there are no additional contour pairs, then step 512 translates hints
20 from the source character to the target character. An example of how this task can
21 be accomplished is given below in the section entitled “Hint Translation”.

22
23
24
25

1 **Matching the Outlines**

2 Referring back to Fig. 2, two glyphs 200, 202 are shown and represent a
3 single character—lowercase “a”. Within this document, the terms “glyphs” and
4 “characters” are used interchangably.

5 Glyph 200 constitutes a source character in the Trebuchet font which has
6 already been hinted and from which hints are to be transferred. Glyph 202
7 constitutes a target character in the Frutiger that has not been hinted and to which
8 hints from glyph 200 are to be transferred. In this particular example, each glyph
9 has two outlines or contours that are to be matched. The outlines for source
10 character 200 are shown at 200a, 200b and the outlines for target character 202 are
11 shown at 202a, 202b.

12 A primary goal of the process described below is to translate hints that refer
13 to control points on the source character or outline to hints that refer to control
14 points on the target character or outline. As indicated above, control points for the
15 two characters are indicated as either a solid dot (to indicate on-curve control
16 points or knots) or an open dot (to indicate off-curve control points or knots).

17 The algorithmic approach described below attempts to match up explicit
18 on-curve knots using features such as contour direction and the presence of
19 extrema. The on-curve knots typically have far more significance to the shape and
20 extents of the contour. Once a match is computed between the on-curve knots, an
21 attempt is made to pair up the remaining knots by simply counting the number of
22 off-curve knots between each pair of matched on-curve knots. If the numbers are
23 equal, then the off-curve knots are paired based solely on their order. Only a very
24 small fraction of hints involve these off-curve knots, but this is done to preserve as
25 many of the source hints as possible.

1 Many glyphs are defined by multiple contours, but there are no restrictions
2 on what order the contours are listed in. Therefore, a first task is to determine
3 which contour goes with which in the two glyphs. In the illustrated example, this
4 is done by enumerating all the possibilities for a one-to-one pairing of the
5 contours. (The hinter rejects input outline pairs with differing numbers of
6 contours). A score is calculated for each pairing as follows. Suppose that the
7 target character is scaled and translated so that its bounding box is equal to that of
8 the source character. For each individual contour within the characters, a sum is
9 computed of the absolute values of the differences between corresponding sides of
10 the contour bounding boxes. This value, summed over all the contours, gives the
11 score for the match, with the lowest value being the best match. Of course, other
12 methods could be used without departing from the spirit and scope of the claimed
13 subject matter.

14

15 **Knot Matching**

16 After the outlines or contours have been matched, an attempt is made to
17 match the knots or control points on each of the outlines so that a basis for
18 transferring hints can be established.

19 In the illustrated example, this is done by pairing individual control points
20 on the outlines, and evaluating the point pairs in accordance with established
21 criteria to ascertain a score that rates their match. Although any suitable criteria
22 can be used, in the illustrated example the criteria take the following form, each of
23 which is explained in more detail below: (1) incoming and outgoing direction, (2)
24 local minimum or maximum, (3) incoming and outgoing lines straight or curved,
25 and (4) band matches.

1 The pairing of the individual control points on the outlines defines a single
2 set of multiple control point pairs. The process is repeated multiple times so that
3 multiple sets are defined. A score is calculated for each set and a set that has the
4 best overall score is selected as a basis for transferring the hints.

5

6 Incoming/Outgoing Direction

7 Fig. 6 shows the characters of Fig. 2 highlighted in a manner that illustrates
8 a first feature that is used for ascertaining a score for a match between control
9 points—incoming/outgoing direction. Specifically, each knot or control point has
10 an incoming and an outgoing direction associated with it. The direction is based
11 on the tangents of the curves touching that knot or control point.

12 In the illustrated example, a clockwise convention is used and the directions
13 are indicated by the arrows that are either pointed toward or away from any one
14 knot. The direction is quantized to one of eight possibilities, corresponding to
15 the eight compass directions, e.g. north, northeast, east, southeast, south,
16 southwest, west, and northwest. A pair of knots on the two characters is assigned
17 from 200 to -200 points based on the similarity of each direction. For example, a
18 knot with an incoming direction of “north,” gets 200 points when matched with
19 another “north” knot, 100 points for a “northeast” or “northwest” match, 0 points
20 for “east” or “west”, -100 points for “southeast” or “southwest”, and -200 points
21 for matching a “south” knot. This score is calculated for both the incoming and
22 outgoing direction.

1 Local Minimum or Maximum

2 Fig. 7 shows the characters of Fig. 2 highlighted in a manner that illustrates
3 a second feature that is used for ascertaining a score for a match between control
4 points—local minimum or maximum. Specifically, each knot can be flagged as a
5 local minimum or maximum in each of the x or y directions. A knot with one of
6 these flags will contribute 150 points when matched with a knot with the same
7 flag, or -150 points when matched to the opposite flag. In the Fig. 7 illustration,
8 local extrema are indicated by the triangles at each of the appropriate knots. A
9 knot may not be an extremum at all in a given direction, in which case any match
10 will not produce a score for this category.

11 Incoming and Outgoing Lines Straight or Curved

12 Fig. 8 shows the characters of Fig. 2 highlighted in a manner that illustrates
13 a third feature that is used for ascertaining a score for a match between control
14 points—incoming and outgoing lines straight or curved. Specifically, each knot
15 has a flag to indicate whether the incoming and outgoing lines are straight (within
16 some tolerance) or curved. In the illustrated example, curved incoming or
17 outgoing lines are indicated as solid lines, while straight incoming or outgoing
18 lines are indicated as dashed lines. Matching these flags produces a score of 100
19 points, but not matching them produces no penalty.

1 **Band Matches**

2 Another feature that is utilized to ascertain a score for a match between
3 control points is referred to as “band matches”. Here, the vertical and horizontal
4 extents of a particular contour are divided into a number of equally spaced bands,
5 e.g. five bands. Corresponding knots on the different contours are given a score of
6 400 points if they fall into the same band, and 200 points if they fall into an
7 adjacent band. Any other distribution does nothing to the score. This process is
8 carried out twice for each contour—once for the horizontal direction and once for
9 the vertical direction. Consider again Fig. 8. There, two different band groupings
10 800, 802 are shown that might be used to assess the horizontal extents of each
11 contour. Accordingly, this feature assesses the desirability of a match by
12 considering whether each control point of a pair falls into one of a plurality of
13 common bands that are defined for each character.

14

15 **Generating the Matches**

16 Fig. 9 shows the characters of Fig. 2 and will assist in understanding an
17 exemplary pairing and matching process in accordance with the described
18 embodiment.

19 To generate knot pairs (control point pairs), an arbitrary starting knot on
20 each contour to be matched is selected and then paired together. For example,
21 consider that the outer contour for each character 200, 202 is undergoing a knot-
22 pairing process. The starting knot on the outer contour for character 200 might be
23 knot 900, while the starting knot on the outer contour for character 202 might be
24 900a. These two knots constitute one pair of knots.

1 After the starting knots are selected, the other knots on the source and target
2 outline are paired in a predetermined manner. This defines multiple knot pairs that
3 collectively define one set of knot pairs. In the illustrated example, the pairing
4 takes place by moving, in a defined direction, around the source contour to the
5 next knot, and pairing each knot with the knot on the target contour whose
6 fractional arc length relative to the starting knot is closest to that of the source
7 knot. In the illustrated example, this next pairing might constitute knots 902,
8 902a. This process is repeated for all of the knots on the source contour until all
9 source knots are paired up with a corresponding target knot. This pairing of all of
10 the source knots generates the set of knot pairs.

11 A local feature score is calculated for each pair of knots, e.g. a local feature
12 score is calculated for pairs 900-900a, 902-902a, etc. One example of how this
13 can be done is given above. Once all of the local feature scores are calculated for
14 all of the knot pairs, the scores are summed to provide an overall score that rates
15 the quality of the overall match of the contours.

16 In the present example, it will be appreciated that the selected starting
17 points on the source and target outline (i.e. points 900, 900a) would not generate a
18 score that reflects a high quality match. This is because point 900 is located on the
19 bottom portion of the character and point 900a is located on the top portion of the
20 character and thus, these points do not have much in common. Thus, these
21 particular pairings of source and target knots is not the most optimal.

22 A different set of knot pairings is now generated and evaluated as described
23 above. In the illustrated example, this can be done by maintaining the starting
24 knot on the source contour (i.e. knot 900) and moving to a different starting knot
25 on the target contour (e.g. knot 902a). The pairing process for each knot on the

1 source contour is repeated as described above so that all of the source knots are
2 paired. For each pair of knots, a local feature score is calculated, and all of the
3 local feature scores are summed to rate the quality of the overall match.

4 This process is repeated so that every knot on the target contour is used as a
5 starting knot for purposes of computing a local feature score.

6 The output of this process provides multiple sets of knot pairs and each set
7 has a score that is calculated as described above. In the illustrated example, the
8 number of scores will be equal to the number of knots on the target outline
9 because each target knot is used as a starting knot for a set of knot pairings.

10 The set of scores is reduced by selecting the five matches with the highest
11 summed local-feature scores. Each of the scores of the reduced set of scores is
12 then further processed to attempt to improve it. This is done by manipulating the
13 individual knot pairings of the set that generated a particular score. Specifically,
14 pairings between knot pairs that generated a negative local-feature score are
15 removed so that the knots are not paired together. Then, unpaired source knots are
16 attempted to be paired with other target knots. Additionally, existing knot pairs are
17 shifted to adjacent target knots. All of this processing is subject to the constraint
18 that the knot match respect the ordering of knots around the contour: e.g., if knot B
19 follows knot A in the source contour, then the partner of knot B should not come
20 before the partner of knot A on the target contour.

21 Once this local improvement processing has been conducted on each of the
22 five top matches, the match with the highest final score is selected as the final
23 match.

1 Fig. 10 shows the results of this matching algorithm. These heuristics work
2 well for a wide variety of character styles, including roman, bold, and italic
3 characters. A matching for a more complex pair of glyphs is shown in Fig. 11.

4 Fig. 12 is a flow diagram that describes steps in a matching algorithm in
5 accordance with the described embodiment. The method is advantageously
6 implemented in software. Step 1200 selects starting knots on the source and target
7 contours. Step 1202 pairs each source knot with a target knot, starting at the
8 starting knots and working around the source contour in a predetermined manner.
9 Step 1204 calculates a local feature score for each knot pair. One example of how
10 this can be done is given above. When all of the local feature scores have been
11 calculated for the individual knot pairs, step 1206 sums the local feature scores to
12 provide an overall score that gives an indication of the quality of the overall
13 match. Advantageously, this process is carried out multiple times so that a set of
14 overall scores is generated for different sets of knot pairings.

15 In the illustrated and described embodiment, each of the target knots is used
16 as a starting knot for a set of knot pairings. Accordingly, step 1208 ascertains
17 whether all of the target knots have been used as a starting knot. If all of the target
18 knots have not been used as a starting knot, step 1210 selects a different target
19 knot as a starting knot (while maintaining the original source starting knot) and
20 returns to step 1202. The result of this step is that a different pair of knots is
21 defined that can generate different local feature scores (step 1204) and a different
22 overall score (step 1206).

23 Once all of the target knots have been used as a starting knot, step 1212
24 selects a predetermined number of the best scores (e.g. the five best scores) and
25 step 1214 attempts to improve each selected score by perturbing knot pairings that

1 contributed to that overall score. Once the highest score has been found, step 1216
2 selects the match with the highest score and uses that match as the basis for
3 transferring hints to the target character.

4 The above description illustrates but one way of implementing a knot-
5 matching process. It is to be understood that other approaches can be
6 implemented without departing from the spirit and scope of the claimed subject
7 matter.

8

9 **Hint Translation**

10 Having produced a “best match” between the knots of the source character
11 outline and those of the target character outline, the hints that are associated with
12 the source character can now be translated or transferred from the source character
13 to the target character.

14 Fig. 13 is a flow diagram that describes an exemplary hint transferring
15 process in which VTT Talk hints are transferred from a source character to a target
16 character. It will be appreciated and understood that while the example is given in
17 the context of VTT Talk hints for TrueType hints, it is possible for the
18 methodologies described below to be implemented in connection with other
19 differently defined hints.

20 Step 1300 defines a match between a source character and a target
21 character. One exemplary way of accomplishing this task is described above in
22 connection with Fig. 12. Once the match is defined, step 1302 transfers hints from
23 the source character to the target character. To transfer the hints in this example,
24 the source font’s VTT Talk hints are parsed (step 1304) and copied (step 1306) to
25 the target font. Knot numbers are replaced (step 1308) as appropriate according to

1 the match. If there is not a match for a knot referenced in a particular statement,
2 the source statement is copied unchanged, but is commented out (step 1310) to
3 mark it as a place that may need special attention by a person reviewing the font.

4 In this manner, hints that were specifically defined for the source character
5 in terms of the source character's knots or control points, are automatically
6 transferred or translated to the target character where the transferred hints now
7 refer specifically to the target character's knots or control points.

9 **CVT Translation**

10 The Control Value Table (CVT) is a central feature of the TrueType hinting
11 mechanism, and no TrueType autohinting scheme would be complete without
12 addressing it.

13 In VTT Talk, entries of the CVT are used via statements such as:

14
15 YLink(14,0,87)
16

17 This statement says, in effect, "move knot 0 up or down so that its vertical
18 distance from knot 14 is equal to the value in CVT entry 87."

19 In the illustrated example, a matching module translates the references to
20 specific knots to their analogues in the new font based upon the matches. In
21 addition, the same CVT entry numbers (e.g. "87") as were used in the original font
22 can still be used, although the values associated with those entry numbers will
23 likely need to be changed. The reason for this is that the old values in the CVT
24 represent distances measured in the source font, which may bear little or no
25 relation to distances in the target font.

1 A solution for which values to insert into the CVT comes from recognizing
2 that the major reason the CVT is used is to take a set of distances that are
3 approximately the same in the outline, and force them to be exactly the same
4 number of pixels in a rendered bitmap. Since the goal is to provide this
5 consistency while changing the outlines as little as possible, the CVT entry will
6 generally contain some average value that is close to all the distances it is going to
7 be used to constrain. Using this knowledge, we can look at all of the uses of a
8 particular CVT entry to estimate what its value should be.

9 As an example, consider Fig. 14 and how this works on the ‘a’ character.
10 The top row shows characters from the font Trebuchet. The typographer hinting
11 this font chose to use CVT entry 87 to represent the height of round, black features
12 in lowercase characters, as indicated by the lines between the control points. Most
13 of the lowercase letters that have round parts reference CVT 87, as indicated in the
14 top row of Fig. 14. The ‘a’ glyph alone uses entry 87 six times—that is, there are
15 six pairs of knots in the ‘a’ whose distance is constrained by CVT entry 87. The
16 bottom row shows the font Frutiger, along with the uses of CVT entry 87 as
17 transferred from Trebuchet by a suitably programmed autohinter. Lines 1400
18 indicate where hints were automatically discarded because the natural distance
19 between these points was too different from the value in the CVT table.
20 Accordingly, hints are disregarded where they appear to be inappropriate for a
21 character of the font to be hinted.

22 The table immediately below shows the “natural” distances between each
23 of these pairs in the Frutiger font for each of the indicated glyphs.
24
25

1 2 3 4 5 6 7 8 9	Glyph	References to CVT entry 87
	"a"	75* 143 143 156 156 164
"b"		111* 113* 156 156
"c"		156 156 160 172
"d"		111* 113* 156 156
"e"		156 156 193*
"f"		155 156
"g"		45* 45* 111* 156 156 178
"h"		156

As indicated, one pair of points in the 'a' is 75 units apart vertically in the unhinted outline, another is 143 units, and so on. To determine the overall value to place in the CVT entry, the median of all these individual values is computed, which in this case is 156 units. The starred numbers in the listing indicate those uses of the CVT entry where the natural outline distance differs by more than 20% from the median value. These values are designated as outliers, and the hints corresponding to those points are removed (i.e. commented out) during the translation process, as they usually represent cases where the shape of the target character differs enough from that of the source character that the CVT constraint is inappropriate. These commented-out constraints correspond to the lines designated at 1400. The commented-out hints can be maintained, however, in commented-out form so that a person fine-tuning the results can quickly see where the autohinter disregarded various hints, thereby possibly necessitating manual work in the area.

1 Note that Trebuchet has a so-called spectacle g, while Frutiger has a multi-
2 story g. In this case, it is likely that two the forms of the 'g' require entirely
3 different hinting strategies, since many of the hints of the source 'g' are simply not
4 appropriate for the target character shape. These inappropriate hints are
5 automatically discarded by the outlier mechanism.

6 It will be appreciated that the CTV has a hierarchical aspect in the sense
7 that entries in the table can refer to other entries whose associated values might be
8 used in certain circumstances. For example, one CTV entry might specify a width
9 value to use above a certain point size, but for point sizes below the certain point
10 size, a different CTV entry is to be used to constrain the width. The above
11 translation of CTV values takes into account these dependencies when the table
12 values are changed.

13

14 **Conditional Hint Specification**

15 In the above discussion, the hints that are being transferred have a
16 characteristic in that they are essentially maintained through all sizes of fonts.
17 There are other hints, however, that are specific to a given size or range of sizes.
18 For example, a typographer may have physically altered the hints of a character at
19 a certain size to make it look better. That is, at a certain size a certain condition
20 may occur that has an undesirable effect on the rendered character, e.g. the
21 character may have a portion that closes up undesirably. In this embodiment, one
22 or more conditions can be defined that result in a certain hint being applied to the
23 character if the condition is satisfied. Advantageously, these conditionally-
24 specified hints can be transferred from a source character to a target character or
25 from a source font to a target font. Thus, if the same condition occurs in the

1 automatically hinted target font, the conditionally-specified hint can be applied in
2 the target character to eliminate the undesirable condition.

3

4 **Effective Removal of Control Points**

5 There are circumstances where certain characters have subtleties in the
6 outlines that look good at higher resolutions, but do not look good at lower
7 resolutions. To deal with this situation, there is a way to manipulate the control
8 points of the character so that the subject control point(s) that give rise to the
9 subtlety are effectively removed. This is done by programmatically relocating the
10 control point so that it lies on top of another control point. By doing this, the
11 control point is effectively, but not actually, removed. Advantageously, in the
12 described embodiment, these effectively removed control points are transferred to
13 the target character so that similar undesirable subtleties in the target character are
14 effectively eliminated as well.

15

16 **ClearType**

17 The above described methods are applicable to ClearType hinting.
18 ClearType is a method of rendering fonts that takes advantage of special properties
19 of LCD screens (such as laptops or PDAs). The TrueType rendering mechanism is
20 unchanged, but for best results, the hinter should alter his or her strategy for
21 hinting the character. Generally speaking, due to increased effective horizontal
22 resolution that ClearType gives, fonts to be displayed with ClearType need much
23 less hinting in the x-direction than fonts intended for CRT display. The above-
24 describe hinting strategy is employable as for TrueType fonts that employ a
25 ClearType strategy.

1

2 Results

3

4 In an example application two TrueType fonts were taken as input: a source
5 font, from which the hints are transferred; and a target font, which is automatically
6 hinted by an application program executing on a computerized hinting system.
7 The application program took under a minute to match the outlines, translate the
8 hints, and create the new CVT for a 256-character font.

9 Once the target font is hinted, however, it should still be reviewed by hand
10 and corrected by an experienced typographer. Even minor errors in the translated
11 hints or CVT can take a considerable amount of time to identify and correct, so the
12 translation should be as accurate as possible in order to be useful.

13 For the explanation that follows, reference is made to Figs. 15-19, each of
14 which visually demonstrates hints that are transferred from a fully hinted Georgia
15 Roman font to an unhinted different font. For example, Fig. 15 shows hints
16 transferred from Georgia Roman to Georgia Bold; Fig. 16 shows hints transferred
17 from Georgia Roman to Bodoni; Fig. 17 shows hints transferred from Georgia
18 Roman to Calisto; Fig. 18 shows hints transferred from Georgia Roman to
19 Perpetua; Fig. 19 shows hints transferred from Georgia Roman to Revival.

20 Figs. 20 and 21 compare the unhinted versions of Sylfaen Sans Bold and
21 Georgia Bold, respectively, to the versions hinted automatically, at 16, 17, and 19
22 ppem, the most commonly used on-screen sizes. In these examples it is clear that
23 most of the objectionable artifacts in the unhinted versions have already been
24 corrected by the automatic hinting. Note, for instance, the improved 'O' shapes
25 and the much more uniform stem weights in both fonts. Still, the autohinted
versions are not perfect; note for instance where the bowl of the Georgia Bold 'b'

1 has narrowed unacceptably, especially at lower sizes. Imperfections like these will
2 need to be corrected by hand.

3 The method described above was evaluated by using the application
4 program to transfer hints between three pairs of fonts within the same family
5 (Sylfaen Sans Bold from Sylfaen Sans, Georgia Bold from Georgia, and Georgia
6 Bold Italic from Georgia Italic) as well as four target fonts from a source font of a
7 different font family (Bodoni, Calisto, Perpetua, and Revival—all from Georgia).
8 The table immediately below summarizes the results of these tests.

Source Font	Target Font	Success rate (%)	Review and clean up (min.)	Manual Hinting (min.)	Savings (%)
Sylfaen Sans	Sylfaen Sans Bold	84%	5.9	9.4	37%
Georgia Italic	Georgia Bold Italic	86%	6.7	7.9	15%
Georgia Roman	Georgia Bold	93%	4.6	7.1	35%
Georgia Roman	Bodoni	78%	3.3	3.3	0%
Georgia Roman	Calisto	74%	3.0	4.3	30%
Georgia Roman	Perpetua	76%	1.2	2.7	56%
Georgia Roman	Revival	82%	1.3	2.3	43%

20
21 In each case, just the alphanumeric glyphs were hinted. The “success rate”
22 column gives the percentage of 62 glyphs in which the transferred hints basically
23 worked. More specifically, for a “successful” glyph, the overall appearance of the
24 glyph conformed to the original outline at high sizes (38 ppem and above) without
25 any stretching or distortion, whereas below 38 ppem there might be some cleaning

1 up to do, but no major reshaping or rethinking of the hints. If a glyph did not
2 conform to its original outline at high sizes or required major reshaping at low
3 sizes, then it was considered “unsuccessful.”

4 As can be seen from the table, the hinter had a fairly high success rate by
5 this measure, especially when hinting characters within the same font family. The
6 next column (Review and cleanup) gives an estimate of the number of minutes
7 required for an experienced typographer to review the results of the autohinter and
8 clean up any problems in the transferred hints. The figures in this column were
9 estimated by performing this process on some 3 to 11 representative glyphs in the
10 target font. These same glyphs were also manually hinted by the same
11 typographer and the times required reported in the following column. Finally, the
12 right-most column provides an estimate of the overall time savings provided by
13 the example-based hinter.

14 Note that the very high success rate of the hinter translates into a more
15 moderate overall time savings, since even a perfectly-hinted font requires time to
16 review, and since a few small problems in the hints can be time-consuming to
17 correct. Still, these savings are significant, considering that a full font of 256
18 characters can take on the order of 20–40 hours for a skilled professional to
19 produce.

20

21 Conclusion

22 Methods and systems for automatically hinting TrueType fonts have been
23 described. The inventive approaches use an existing, fully-hinted font as a
24 template, thereby allowing the hints of one font to be transferred to another. This
25 translation process includes estimation of the control value table (CVT) entries

1 that are used to unify feature sizes across a font. Utility is derived from the
2 described matching algorithm from not only its simplicity, but from its ability to
3 be applied to and work well for a wide variety of character shapes, including
4 serifed and italic fonts.

5 An important advantage of the described approach over previous
6 autohinters is that the described approach preserves the hand-crafted hinting
7 strategy, built by a professional typographer, in the newly hinted font. Thus, the
8 translated hints provide an excellent starting point and generally require only
9 minor cleanup and adjustment.

10 Although the invention has been described in language specific to structural
11 features and/or methodological steps, it is to be understood that the invention
12 defined in the appended claims is not necessarily limited to the specific features or
13 steps described. Rather, the specific features and steps are disclosed as preferred
14 forms of implementing the claimed invention.

15

16

17

18

19

20

21

22

23

24

25